Streaming Readout for EIC Detectors

Douglas Hasell

for the

Streaming Readout Consortium

EIC Detector R&D Meeting
Catholic University of America
July 26, 2018

Streaming Readout Consortium

S. Ali, V. Berdnikov, T. Horn, I. Pegg, R. Trotta
Catholic University of America, Washington DC, USA

M. Battaglieri (Co-PI), A. Celentano
INFN, Genova, Italy

J.C. Bernauer† (Co-PI), D.K. Hasell, R. Milner
Massachusetts Institute of Technology, Cambridge, MA

C. Cuevas, M. Diefenthaler, R. Ent, G. Heyes, B. Raydo, R. Yoshida
Thomas Jefferson National Accelerator Facility, Newport News, VA

New members welcome !

Yulia Furletova (JLab) and GEM-TRD/T group will collaborate.

† at Stony Brook University beginning September, 2018

NAS Report on EIC Science - 7/24/2018

"*The continued rapid pace of technological development is starting to enable a transition from the event-oriented and triggered data-acquisitions of past and current experiments in nuclear and high-energy physics to data models where detector subsystems deliver time-stamped streams of data for processing with increasingly integrated and advanced computing resources in real time.*"

Streaming readout will greatly benefit future EIC experiments

- luminosity 10^{33} – 10^{35} cm^{-2}·s^{-1}
- event rates 100 – 1,000 $\times$ rates at HERA
- detectors with $\mathcal{O}(10^6)$ channels

Streaming readout impacts detector, electronics, software, and analysis

We must begin working to understand this now

Ideal Readout Scheme - Schematic for One Channel

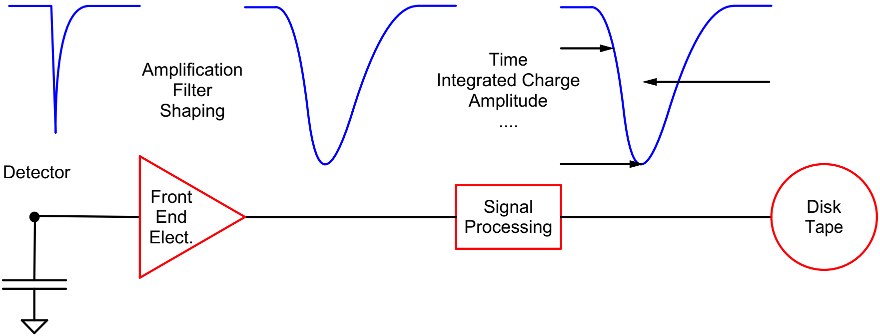

Consider this simple readout scheme

- detector channel experiences an event and generates a signal
- front-end electronics amplifies, filters, and shapes appropriately
- signal processing extracts time, integrated charge, amplitude, . . .
- results are written to disk or tape for offline analysis

Ideal Readout Scheme - Schematic for One Channel

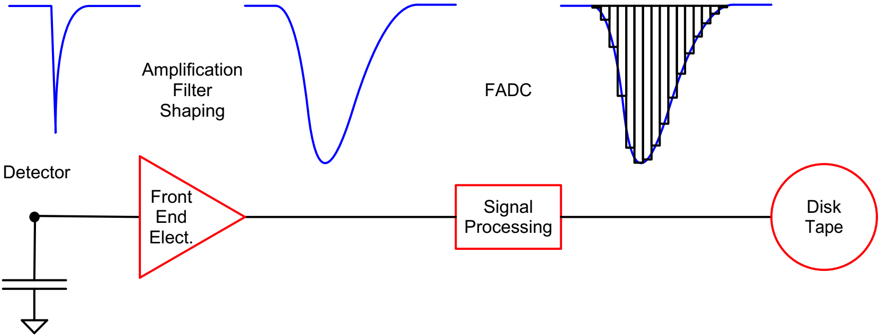

More modern scheme might use a flash ADC or TDC

- samples the signal at some rate and writes the results
- signal processing possible offline
- time, charge, amplitude, etc. parameters extracted as needed
- more information also available: pile-up, signal shape, noise, . . .

Ideal Readout Scheme - for More than One Channel

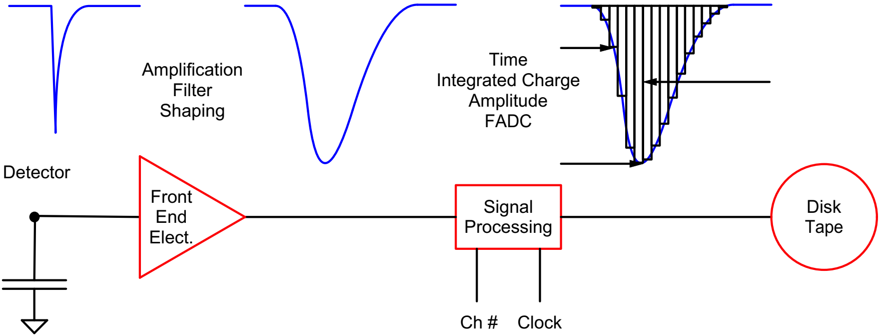

For more than one channel need to add channel ID and time stamp

Scheme provides trigger free, independent readout of all channels

- BUT ! ⇒ breaks down with increasing event size × event rate
- ⇒ limited by computer resources, time, manpower, and $$$
- need to discard noise, background, and unwanted events

Traditional Triggered Readout Scheme

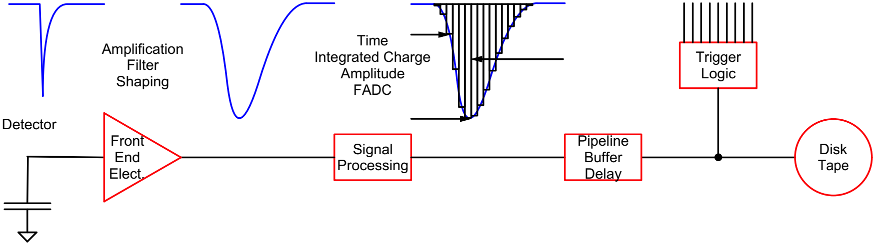

Traditional solution is a triggered readout scheme

- signal or signal parameters stored in pipeline / buffer or delayed
- trigger decision based on fast signals from subset of detector channels
- often special detectors, hardware, and electronics involved
- sometimes multiple levels of buffers and triggers required

Trigger logic can veto or select events but ...

- decision based on a subset of the data
- logic based on past experience and expectations
- finds what is expected → might miss the unexpected

Streaming Readout is Possible

For the EIC we need to start working towards this goal now !

- implications for the detectors, electronics, software, and analysis
- the whole chain needs coordination and standardisation

Leverage advances and falling costs of electronics, computers, storage, ...

- ASICs → multiplexed ADC/TDC chips, rad. hard, low power, ...
- FPGAs → affordable, multi-channel, digital signal processing
  - now with UNIX OS to simplified programming
- high bandwidth copper and optical fibre networking solutions
- affordable, multi-core CPU clusters to analyse data in real time
- reconstruction algorithms: neural networks, machine learning, ...
- TPU chips - artificial intelligence accelerator ASIC

Many experiments already moving toward streaming readout !

Readout Scheme at LHCb

JINST 8 (2013) P04022

Real-Time Processing

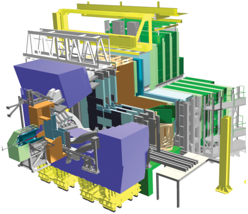

Simple feature-building, e.g. in FPGAs, required to reduce the data rate.*

1 TB/s post zero suppression

50 GB/s

1 MHz

Heavy use of machine learning: V.Gligorov, MW, JINST 8 (2012) P02013.

Real-time reconstruction for all charged particles with $p_T > 0.5$GeV.

8 GB/s

Data buffered on 10 PB of disk.

Real-time calibration & alignment.

Full real-time reconstruction for all particles available to select events.

0.7 GB/s (mix of full + partial events)

*LHCb will move to a triggerless-readout system for LHC Run 3 (2021-2023), and process 5 TB/s in real time on the CPU farm.

Slide provided by Mike Williams (MIT)

Streaming Readout Scheme

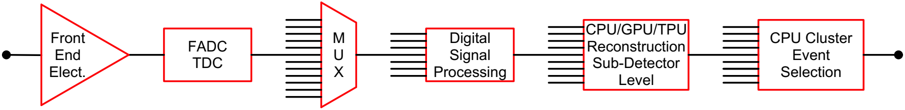

Goal: Do as much as possible on-line and save everything that is useful

- extract relevant parameters for all channels
- extract time, position, momenta, angles, energy, . . .
- discard noise and background hits when possible
- save high-level information, not hits !
- combine data in sub-detector or sectors to form event segments
- option to organised into events or not !
- all detector information available for making decisions

Consequences for detectors, electronics, software, and analysis !

ASIC Board on Detector

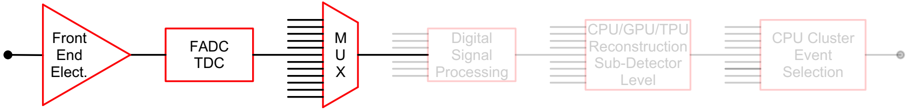

Specifically designed boards mounted on the detector

- FEE must be matched to detector requirements
- multi-channel ASIC chips available with FADCs and multiplexing
    - *e.g.* 64 ch, 12 bit, 1 GSPS, < 20 mW/ch, rad. hard, $< \$10$/ch
- instead of FADC chip TDC chip could be used if appropriate
- zero-suppression, only viable signals passed on
- copper (supply power) or optical fibre (electrical isolation) to DSP

FPGA Signal Processing

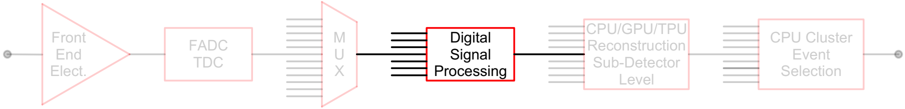

Digital signal processing (DSP) with inexpensive FPGA boards

- can be situated at some distance (accessible) from the detector
- multiple input channels, de-multiplex FADC information
- analyse input channels in parallel for: time, charge, amplitude, ...
  - compress data from high frequency samples to a few parameters
  - flexibility to optimise software as needed
  - data monitoring in parallel for slow control, calibration, ...
  - add channel number and time stamp to channel data stream
- add channel ID and time stamp for each channel's data stream
- high bandwidth output to CPU / GPU / TPU reconstruction

Time synchronisation over all channels important !

Reconstruction at Sub-Detector Level

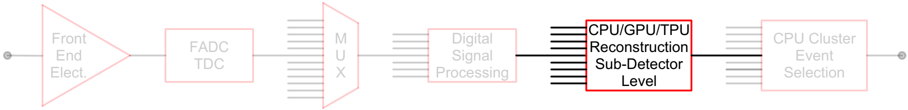

CPU / GPU / TPU analyses sub-detector information

- analysis / reconstruction of localised data
- form clusters, track segments, PID, ...
- high-level parameters: position, time, momenta, angle, energy, ...
- save this high-level information in time slots
- discard "raw" data $\rightarrow$ event size becomes manageable

Perhaps this is sufficient ? Output at this stage ?

- no need to form complete events, leave for off-line analysis
- next stage, organisation into events optional

Combine Sub-Detector Information for Event Selection

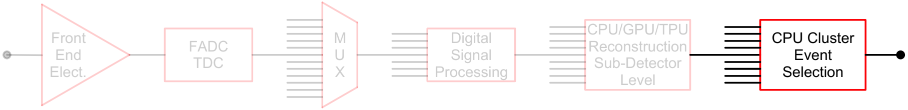

Combine sub-detector data to form complete tracks / events

- connect track segments, associate with calorimeter clusters, etc.
- categorise events
- write to disk for off-line physics analysis
- directly produce DSTs $\rightarrow$ reduce event size

Events immediately ready for physics analysis

- reduce time to achieve results / publication

Streaming Readout Issues

Readout scheme and detector, software, and analysis strongly coupled

- FEE needs to be tailored to detector and its requirements
- FPGA signal processing software specific to detector requirements
- possible to identify basic designs that can be modified to suit
    - change shaping time, sampling rate, number of bits, *etc.* as needed
    - code libraries to extract timing, amplitude, charge, analysis, *etc.*
- standardise components to be adaptable to many experiments
    - plug-n-play for streaming readout like old CAMAC and NIM modules

Detector must also be designed with streaming readout in mind

- detector can not require a trigger to initiate readout
- prefer fast detector response, avoid long response times
- plan ahead to simplify analysis
- detector, readout, and analysis must be designed together

Streaming Readout Development

MIT workshops to review status and plan for future

- first workshop January 27, 2017, second January 29–30, 2018
- now holding monthly video meetings
- electronics groups from JLab, BNL, ANL, SLAC involved
- TOPSIDE, sPHENIX, SOLid, LHCb, GlueX, DarkLight experimenters
- representatives from industry CAEN, AlphaCore

Lots of groups working towards streaming readout

- national labs have great resources but “doing their own thing”
- ASIC chip developers very interested but need direction
- need to coordinate efforts

Streaming Readout R&D Proposal

Aim to identify and quantify parameters for streaming readout by studying its implementation on various prototype systems.

Deliverable is a report detailing:

- relevant aspects about each test prototype's performance,
- definition of streaming readout parameters relevant to EIC,
- list of parameters for various detector technologies,
- estimates of parameters based on current technologies, and
- extrapolation of parameters to detector construction time frame.

Detectors, electronics, software, and analysis interdependent

Must begin work on streaming readout now !

Plan for Work at JLab

Streaming readout studies using the TDIS GEM readout for the rTPC

- use rTPC test stand being assembled at JLab
- components and software based on the ALICE DAQ upgrade
- gain experience to guide design of streaming readout

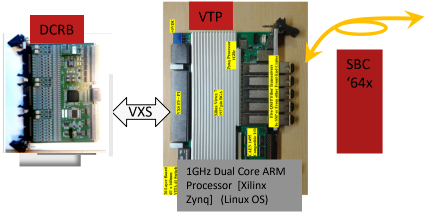

Crate-less streaming readout

- alternative to standard VME
- use VXS
- 250 MHz FADC and DCRB
- VXS trigger processor 20 Gb
- 4GB buffering

No funding request for this part of the proposal

Plan for Work at MIT

Work on design of on-detector electronics

- front-end electronics
- ASIC chips in collaboration with Alphacore
  - preamplifiers, charge sensitive, shaper
  - 300-1300$\times$ gain for 0-300 pF
  - FADC 10-12 bit, 100-200 MSPS
- output matched to FPGA test boards
- various prototype detectors SiPM
  - SiPM connected to scintillating strip
  - GEM detectors
  - $PbWO_4$ with SiPM, APD or PMT together with INFN Genova

No funding request for this part of the proposal

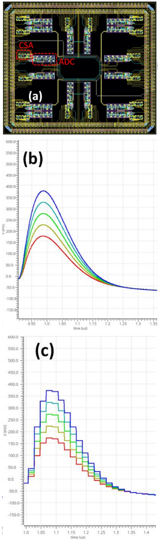

Plan for Work at INFN Genova and CUA

Streaming readout scheme for electromagnetic calorimetry

- investigate technologies suitable for various calorimeters
- SiPM, APD, PMT readout
- $PbWO_4$ calorimeters

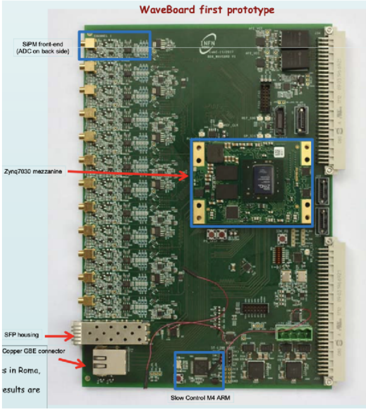

Digitiser board at INFN-Genova

- 12 channels
- 250 MHz, 14 bit
- FEE matched to SiPM readout
- Zynq7030 FPGA dual core ARM
- copper and fiber output
- VME for power only

15 k$ requested both CUA and INFN Genova

Plan for Work at Stony Brook University

Propose to investigate data transport architectures for streaming readout

- study different protocols: IP, TCP, UDP, Ethernet, ...
- develop application layer solutions for data, slow control, ...
- framework for online data processing
- study architectures for CPU clusters to best use resources

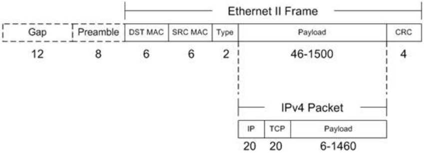

20 k$ requested for this part of the proposal

Funding Request for Streaming Readout

Institute	Equipment	Travel	Sum
Catholic University of America	10 k$	5 k$	15 k$
INFN Genova	10 k$	5 k$	15 k$
Stony Brook University	15 k$	5 k$	20 k$
MIT	0 k$	0 k$	0 k$
JLab	0 k$	0 k$	0 k$
Total	35 k$	15 k$	50 k$

Conclusion

Streaming readout is possible !

- advances in technology → ASIC, FPGA, muti-core CPU, ...
- falling costs in electronics, computing, storage, and networking
- advances in software → neural networks, machine learning, TPUs, ...
- can expect further improvements over the next decade

Future EIC experiments will benefit from streaming readout

Implications for detector, electronics, software, and analysis

Important that streaming readout approach be endorsed now

So we can begin work !

Thank You